

CLAIMS

What is claimed is:

1 1. A switching device comprising
2 a segmentation unit to receive data packets of variable size and to store
3 the data packets as segments, wherein the data packets received having a length
4 greater than a maximum segment length are divided into multiple segments;
5 a plurality of queues to store the segments, wherein the queues are associated with
6 destinations;
7 a scheduler to generate a schedule including a data path from at least one queue
8 to an associated destination;
9 a framer to aggregate a plurality of segments to form a frame, wherein the frame
10 has a maximum frame length, wherein the segments are retrieved from the at least one
11 queue, and wherein the frame may contain segments associated with different data
12 packets; and
13 a transmitter to transmit the frame to the associated destination, wherein the
14 segments within the frame are transmitted together.

1 2. The device of claim 1, wherein said segmentation unit divides the data
2 packets having a length greater than the maximum segment length to form at least a first
3 segment having the maximum segment length.

1 3. The device of claim 2, wherein a last segment formed by said
2 segmentation unit may be less than the maximum segment length.

1 4. The device of claim 1, wherein said segmentation unit identifies a final
2 segment belonging to a particular data packet.

1 5. The device of claim 1, wherein said scheduler generates the schedule
2 based on requests for data packets from particular queues to be transmitted to particular
3 destinations.

1 6. The device of claim 5, wherein said scheduler matches the requests and
2 resolves conflicts in order to generate the schedule.

1 7. The device of claim 1, wherein said plurality of queues include one or
2 more queues per destination based on priorities associated with the destinations.

1 8. The device of claim 1, wherein said framer forms the frame by
2 aggregating segments from some subset of the one or more associated queues.

1 9. The device of claim 1, wherein said framer forms the frame by
2 aggregating complete segments.

1 10. The device of claim 9, wherein said framer pads the frame to reach the
2 maximum frame length.

1 11. The device of claim 1, wherein said transmitter transmits the segments
2 making up a particular data packet in order.

1 12. The device of claim 1, wherein said transmitter interleaves segments
2 associated with different data packets.

1 13. The device of claim 1, further comprising a striper to stripe the frame
2 across a plurality of channels.

1 14. The device of claim 1, further comprising a switch fabric to provide a
2 configurable connection between a plurality of sources and a plurality of destinations,
3 wherein said switch fabric is configured by said scheduler, and wherein said switch fabric
4 transmits frames from at least a subset of the plurality of sources to at least a subset of the
5 plurality of destinations during a schedule cycle.

1 15. A method comprising:
2 receiving data packets of variable size;
3 forming segments from the received data packets, wherein the data packets
4 received having a length greater than a maximum segment length are divided into
5 multiple segments;
6 storing the segments in queues, wherein the queues are associated with
7 destinations;
8 selecting a data path from at least one queue to an associated destination;
9 aggregating a plurality of segments to form a frame, wherein the frame has a
10 maximum frame length, wherein the segments are retrieved from the at least one queue,
11 and wherein the frame may contain segments associated with different data packets;
12 transmitting the frame to the associated destination, wherein the segments within
13 the frame are transmitted together.

1 16. The method of claim 15, wherein said forming segments includes dividing
2 the data packets having a length greater than the maximum segment length into at least a
3 first segment having the maximum segment length, and wherein a last segment formed
4 may be less than the maximum segment size.

1 17. The method of claim 15, wherein said aggregating includes aggregating
2 complete segments.

1 18. The method of claim 15, wherein said transmitting includes striping the
2 frame across a plurality of channels.

1 19. An apparatus comprising
2 a plurality of sources;
3 a plurality of destinations;
4 a receiver to receive frames, wherein the frames include a plurality of segments
5 and the segments may be associated with multiple data packets;
6 a deframer to extract the segments from the frame;
7 a queue to store the segments;
8 a state monitor to track complete packets contained within the queue;
9 a reassembly unit to combine the segments making up a complete packet together
10 to generate the packet;
11 a scheduler to generate a schedule for transmitting frames; and
12 a switch fabric to selectively connect the plurality of sources to the plurality of
13 destinations based on the schedule, wherein said switch fabric transmits frames from at
14 least a subset of the plurality of sources to at least a subset of the plurality of destinations.

1 20. The apparatus of claim 19, wherein the frames are striped over a plurality
2 of channels and said receiver is a deskewer that receives the frame over the plurality of
3 channels and deskews the stripes and combines them into the frame.

1 21. The apparatus of claim 19, further comprising an error checker to check
2 for errors in the frame.

1 22. The apparatus of claim 21, wherein said error checker checks for errors by
2 comparing a grant received to the frame.

1 23. The apparatus of claim 19, wherein said state monitor determines
2 complete packets by detecting an end of packet.

1 24. A method comprising
2 receiving frames of data from a plurality of sources for a plurality of destinations
3 over a switch fabric, wherein a frame includes a plurality of segments and wherein the
4 segments may be associated with multiple data packets;
5 extracting the segments from the frame;
6 storing the segments in a queue;
7 monitoring complete packets contained within the queue; and
8 combining the segments making up a complete packet together to generate the
9 packet.

1 25. The method of claim 24, further comprising checking for errors in the
2 frame.

1 26. The method of claim 24, wherein said monitoring includes detecting an
2 end of packet.

1 27. A store and forward device comprising

2 a plurality of Ethernet cards to receive data from and transmit data to external
3 sources, wherein the plurality of Ethernet cards include
4 an ingress interface module to receive packets of variable size, store the
5 packets as segments in queues associated with at least some subset of destination and
6 priority, aggregate multiple segments together to form a frame, and transmit the segments
7 within the frame together, wherein the segments stored in the queues have a segment
8 maximum size and the packets having a larger size are divided into multiple segments,
9 and wherein the frames have a maximum frame size, include a plurality of segments and
10 may include segments associated with different packets, and
11 an egress interface module to receive frames, divide the frames into
12 segments, store the segments in a queue, monitor the queues for complete packets, and
13 reassemble a packet after the segments making up the packet are received;
14 a switching matrix to provide selective connectivity between said Ethernet cards;
15 a backplane consisting of a plurality of channels to connect said plurality of
16 Ethernet cards to said switching matrix; and
17 a scheduler to select connectivity between Ethernet cards and to configure said
18 switching matrix accordingly.

1 28. The device of claim 27, wherein said backplane also connects said
2 plurality of Ethernet cards to said scheduler.

1 29. The device of claim 28, wherein said backplane uses same channels to
2 connect said Ethernet cards to said switching matrix and to said scheduler.

1 30. The device of claim 28, wherein said backplane provides a logical
2 separation of data path between said Ethernet cards and said switching matrix and
3 scheduling path between said Ethernet cards and said scheduler.

1 31. The device of claim 27, wherein the maximum frame size is selected
2 based on at least some subset of configuration time of data path, scheduling time for
3 scheduler, and complexity of scheduling algorithms.

1 32. The device of claim 27, wherein the backplane is optical.

1 33. The device of claim 27, wherein the switching matrix is optical.

1 34. A store and forward device comprising
2 a plurality of ingress interface modules to receive packets of variable size, store
3 the packets as segments in queues, aggregate multiple segments together to form a frame,
4 and transmit the segments within the frame together, wherein the segments stored in the
5 queues have a segment maximum size and the packets having a larger size are divided
6 into multiple segments, and wherein the frames have a maximum frame size and include
7 a plurality of whole segments and may include segments associated with different
8 packets;

9 a plurality of egress interface modules to receive frames, divide the frames into
10 segments, store the segments in queues, monitor the queues for complete packets, and
11 reassemble a packet after the segments making up the packet are received;

12 a switching matrix to provide selective connectivity between said plurality of
13 ingress interface modules and said plurality of egress interface modules;

14 a backplane consisting of a plurality of channels to connect said plurality of
15 ingress interface modules and said plurality of egress interface modules to said switching
16 matrix; and

17 a scheduler to select connectivity between said plurality of ingress interface
18 modules and said plurality of egress interface modules, and to configure said switching
19 matrix accordingly.

1 35. The device of claim 34, wherein a pipeline schedule is implemented that
2 includes at least some subset of
3 said ingress interface modules forming and transmitting requests to said
4 scheduler,
5 scheduler computing a schedule based on the requests,
6 said scheduler transmitting grants to associated ingress interface modules and
7 configuring said switching matrix, and
8 said ingress interface modules transmitting the frames in response to the grants.

1 36. The device of claim 34, wherein said switching matrix extracts a clock
2 from frames received, re-times the frames to a common internal clock domain, aligns any
3 configuration changes provided by said scheduler to the boundaries of the frames, and
4 maintains an uninterrupted data stream to the egress interface modules by transmitting an
5 idle stream when no frames are being transmitted.